

Dimitrios Tsoumakos, PhD

0.1. Personal

Date of Birth	14/04/1976
Place of Birth	ATHENS
Telephone	+30-210-7722867
E-mail	dtsouma@gmail.com
Web page	http://www.cslab.ntua.gr/~dtsouma/

0.1.1. Studies

2002- 2006	Doctor of Philosophy (PhD), Computer Science Department, University of Maryland, College Park, 20742, USA Title of Thesis: ' <i>Search, Replication and Data Dissemination for Unstructured P2P Networks</i> ' Advisor: Prof. Nicholas Roussopoulos
2000- 2002	Master of Science (MSc), Computer Science Department, University of Maryland, College Park, 20742, USA Title of Thesis: ' <i>Tracking the Locality of change in Web Pages</i> ' Advisor: Prof. Nicholas Roussopoulos
1994- 1999	Diploma in Electrical and Computer Engineering, School of Electrical and Computer Engineering, National Technical University of Athens, Greece Title of Thesis: ' <i>Electronic Study Guide of the ECE Department</i> ' Advisor: Prof. Timos Sellis

0.1.2 Scholarships and Awards

2019	Best Paper Runner up Award in the 31st International Conference on Scientific and Statistical Database Management (SSDBM 2019)
2013	Best Paper Award in the 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2013)
2000-2002	UMIACS (University of Maryland Institute for Advanced Computer Studies) scholarship
1998	«Karydoyiannis» scholarship for high G.P.A. (9.6/10.0)
1994	«Papastavridis» scholarship for my admittance to the ECE department (6th out of 160)

0.1.3 Publications

0.1.3.1 Theses

- [T01] **Tsoumakos, D.** 2006. 'Search, Replication and Data Dissemination for Unstructured P2P Networks', PhD Thesis, Computer Science Department, University of Maryland, College Park, 20742, USA.
- [T02] **Tsoumakos, D.** 2002. 'Tracking the Locality of change in Web Pages', MSc Thesis, Computer Science Department, University of Maryland, College Park, 20742, USA.
- [T03] **Tsoumakos, D.** 1999. 'Electronic Study Guide of the ECE Department', Diploma Thesis, School of Electrical and Computer Engineering, National Technical University of Athens, Greece.

0.1.3.2 Articles in Refereed International Journals

- [J01] **Mytilinis, I., Tsoumakos, D., Koziris, N., 2020.** 'Workload-Aware Wavelet Synopses for Sliding Window Aggregates'. In: *Distrib. Parallel Databases (2020)*. <https://doi.org/10.1007/s10619-020-07307-w>.

- [J02] **Doka, K., Kosmatopoulos, A., Papadopoulos, A., Sioutas, S., Tsihlias, K., Tsoumakos, D., 2020.** 'Dynamic planar range skyline queries in log logarithmic expected time'. In: *Information Processing Letters, Volume 162, 2020, 105990.*
- [J03] **Doka, K., Mytilinis, I., Papailiou, N., Giannakouris, V., Tsoumakos, D., Koziris, N., 2019.** 'Multi-engine Analytics with IReS'. In: *Castellanos M., Chrysanthis P., Pelechrinis K. (eds) Real-Time Business Intelligence and Analytics. BIRTE 2015, BIRTE 2016, BIRTE 2017. Lecture Notes in Business Information Processing, vol 337. Springer, Cham.*
- [J04] **Mytilinis, I., Tsoumakos, D., Koziris, N., 2019.** 'Scaling the Construction of Wavelet Synopses for Maximum Error Metrics'. In *IEEE Transactions on Knowledge and Data Engineering, vol. 31, no. 9, pp. 1794-1808, 1 Sept. 2019.*
- [J05] **Giannakopoulos, I., Konstantinou, I., Tsoumakos, D., Koziris, N., 2018.** 'Cloud application deployment with transient failure recovery', In *Journal of Cloud Computing. Advances, Systems and Applications, 2018, 7:11.*
- [J06] **Nodarakis, N., Pitoura, E., Sioutas, S., Tsakalidis, A., Tsoumakos, D., Tzimas, G., 2016.** 'kdANN+: A Rapid AkNN Classifier for Big Data', In *Transactions on Large-Scale Data- and Knowledge-Centered Systems, 23: 139-168 (2016).*
- [J07] **Xing, W., Jie, W., Tsoumakos, D., Ghanem, M., 2015.** 'A network approach for managing and processing big cancer data', In *Cluster Computing 18(3): 1285-1294 (2015).*
- [J08] **Sioutas, S., Sakkopoulos, E., Panaretos, A., Tsoumakos, D., Gerolymatos, P., Tzimas, G., Manolopoulos, Y., 2014.** 'D-P2P-Sim+: A novel distributed framework for P2P protocols performance testing', In *Journal of Systems and Software, Nov. 2014*
- [J09] **Sioutas, S., Panaretos, A., Karydis, I., Tsoumakos, D., Tzimas, G., Tsohis, D., 2012.** 'SART: Speeding up Query Processing in Sensor Networks with an Autonomous Range Tree Structure', In *SIGAPP Appl. Comput. Rev., Vol. 12, No. 3, 60-74, Sep. 2012*
- [J10] **Asiki, A., Tsoumakos, D., Koziris, N., 2012.** 'A DHT-Based System for the Management of Loosely Structured, Multidimensional Data', In *Transactions on Large-Scale Data- and Knowledge-Centered Systems(TLDKS), 6, 2012, 134-166.*
- [J11] **Doka, K., Tsoumakos, D., Koziris, N., 2011.,** 'Brown Dwarf: A Fully-Distributed, Fault-Tolerant Data Warehousing System', In *Journal of Parallel and Distributed Computing 71(11), 1434-1446, Nov. 2011.*
- [J12] **Konstantinou, I., Tsoumakos, D., Koziris, N., 2011.** 'Fast and Cost-Effective Online Load-Balancing in Distributed Range-Queryable Systems', In *Transactions on Parallel and Distributed Systems, 22(8), 1350-1364, Aug. 2011.*
- [J13] **Doka, K., Tsoumakos, D., Koziris, N., 2010.,** 'Online Querying of d-Dimensional Hierarchies', *Journal of Parallel and Distributed Computing, 71(3), 424-437, March 2011.*
- [J14] **Asiki, A., Tsoumakos, D., Koziris, N., 2010.** 'Distributing and Searching Concept Hierarchies: An Adaptive DHT-based System', *Cluster Computing, 13(3): 257-276.*
- [J15] **Chazapis, A., Asiki, A., Tsoukalas, G., Tsoumakos, D., Koziris, N., 2010.** 'Replica-aware, multi-dimensional range queries in Distributed Hash Tables', *Computer Communications 33(8): 984-996.*
- [J16] **Asiki, A., Doka, K., Zissimos, A., Konstantinou, I., Tsoumakos, D., Koziris, N., Tsanakas, P., 2009.** 'A Grid middleware for data management exploiting Peer-to-Peer techniques', In *FUTURE GENERATION COMPUTER SYSTEMS, 25: 426-435.*
- [J17] **Kantere, V., Tsoumakos, D., Sellis, T., Roussopoulos, N., 2009.** 'GrouPeer: Dynamic Clustering of P2P Databases', *Information Systems 34(1): 62-86*
- [J18] **Kantere, V., Tsoumakos, D., Sellis, T., 2008.** 'A Framework for Semantic Grouping in P2P Databases', *Information Systems 33(7-8): 611-636*
- [J19] **Tsoumakos, D., Roussopoulos N., 2003.** 'Probabilistic Knowledge Discovery and Management for P2P Networks', *P2P Journal, November 2003*

0.1.3.3 Book Chapters

0.1.3.4 Articles in Refereed International Conferences

- [C01] **Giannakopoulos, I., Tsoumakos, D., 2020.** 'Content-based Analytics: Moving Beyond Data Size'. In *Proceedings of the Sixth IEEE International Conference on Big Data Computing Service and Machine Learning Applications, August 2020, Oxford, UK (online).*
- [C02] **Mytilinis, I., Tsoumakos, D., Koziris, N., 2019.** 'Maintaining Wavelet Synopses for Sliding-Window Aggregates'. In *Proceedings of the 31st International Conference on Scientific and Statistical Database Management (SSDBM), July 23-25 2019, Santa Cruz, CA, USA.*

- [C03] **Bakogiannis, T., Giannakopoulos, I., Tsoumakos, D., Koziris, N.**, 2019: 'Apollo: A Dataset Profiling and Operator Modeling System'. In Proceedings of the 2019 ACM SIGMOD/PODS International Conference on Management of Data, June 30 - July 5, 2019 Amsterdam, The Netherlands.
- [C04] **Bakogiannis, T., Giannakopoulos, I., Tsoumakos, D., Koziris, N.**, 2019: 'Predicting Graph Operator Output over Multiple Graphs'. In Proceedings of the 19th International Conference on Web Engineering (ICWE), June 11 - 14, 2019, Daejeon, Korea.
- [C05] **Giannakopoulos, I., Tsoumakos, D., Koziris, N.**, 2018. 'A Content-Based Approach for Modeling Analytics Operators'. In Proceedings of the 2018 Conference on Information and Knowledge Management (CIKM 2018), April 22-26, 2018, Lingotto, Turin, Italy.
- [C06] **Giannakopoulos, I., Tsoumakos, D., Koziris, N.**, 2018. 'Towards an Adaptive, Fully Automated Performance Modeling Methodology for Cloud Applications'. In Proceedings of the IEEE International Conference on Cloud Engineering 2018 (IC2E 2018), April 17-20, 2018, Orlando, FL, USA.
- [C07] **Giannakopoulos, I., Tsoumakos, D., Koziris, N.**, 2017. 'A Decision Tree Based Approach Towards Adaptive Modeling of Big Data Applications'. In Proceedings of the 2017 IEEE International Conference on Big Data (BigData 2017), December 11-14, 2017, Boston, MA, USA.
- [C08] **Singh, A., Tsoumakos, D.**, 2017. 'Towards an Algebraic Cost Model for Graph Operators'. In Proceedings of the 3rd International Workshop on Algorithmic Aspects of Cloud Computing (ALGO CLOUD 2017), in conjunction with the ALGO 2017 Conference, September 4-8, 2017, Vienna, Austria.
- [C09] **Giannakopoulos, I., Konstantinou, I., Tsoumakos, D., Koziris, N.**, 2017. 'AURA: Recovering from Transient Failures in Cloud Deployments' (demo paper). In Proceedings of the 17th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID 2017), May 14-17, 2017, Madrid, Spain.
- [C10] **Doka, K., Papailiou, N., Giannakouris, V., Tsoumakos, D., Koziris, N.**, 2016. 'Mix 'n' Match Multi-Engine Analytics'. In Proceedings of the 2016 IEEE International Conference on Big Data (BigData 2016), December 5-8, 2016, Washington DC, USA.
- [C11] **Giannakouris, V., Papailiou, N., Tsoumakos, D., Koziris, N.**, 2016. 'MuSQLE: Distributed SQL Query Execution Over Multiple Engine Environments'. In Proceedings of the 2016 IEEE International Conference on Big Data (BigData 2016), December 5-8, 2016, Washington DC, USA.
- [C12] **Korasidis, N., Giannakopoulos, I., Doka, K., Tsoumakos, D., Koziris, N.**, 2016. 'Fair, Fast and Frugal Large-Scale Matchmaking for VM Placement'. In Proceedings of the 2nd International Workshop on Algorithmic Aspects of Cloud Computing (ALGO CLOUD 2016), in conjunction with the ALGO 2016 Conference, August 22, 2016, Aarhus, Denmark.
- [C13] **Giannakopoulos, I., Konstantinou, I., Tsoumakos, D., Koziris, N.**, 2016. 'Recovering from Cloud Application Deployment Failures through Re-execution'. In Proceedings of the 2nd International Workshop on Algorithmic Aspects of Cloud Computing (ALGO CLOUD 2016), in conjunction with the ALGO 2016 Conference, August 22, 2016, Aarhus, Denmark.
- [C14] **Papailiou, N., Doka, K., Giannakouris, V., Papaioannou, V., Tsoumakos, D., Koziris, N.**, 2016. 'Robust and Adaptive Multi-Engine Analytics using IReS (demo paper)'. In Proceedings of the 10th International Workshop on Enabling Real-Time Business Intelligence (BIRTE 2016), in conjunction with the VLDB 2016 Conference, New Delhi, India, September 9, 2016.
- [C15] **Mytilinis, I., Tsoumakos, D., Koziris, N.**, 2016. 'Distributed Wavelet Thresholding for Maximum Error Metrics'. In Proceedings of the 2016 ACM SIGMOD/PODS International Conference on Management of Data, June 26 - July 1, 2016, San Francisco, USA.
- [C16] **Doka, K., Filatov, M., Giannakouris, V., Kantere, V., Koziris, N., Mantas, C., Papailiou, N., Papaioannou, V., Tsoumakos, D.**, 2016. 'Optimizing, Planning and Executing Analytics Workflows over Multiple Engines'. In Proceedings of the 1st International Workshop on Multi-Engine Data AnaLytics (MEDAL) (demo paper), in conjunction with the EDBT/ICDT 2016 Joint Conference, Bordeaux, France, March 15, 2016.
- [C17] **Giannakopoulos, I., Karras, P., Tsoumakos, D., Doka, K., Koziris, N.**, 2015. 'An Equitable Solution to the Stable Marriage Problem'. In Proceedings of the 27th IEEE International Conference on Tools with Artificial Intelligence (ICTAI 2015), 9-11 November, Vietri sul Mare, Italy.
- [C18] **K. Doka, K., Xue, M., Tsoumakos, D., Karras, P., Cuzzocrea, A., Koziris, N.**, 2015. 'Heterogeneous k-Anonymization with High Utility'. In Proceedings of the 2nd International Workshop on Privacy and Security of Big Data (in conjunction with the 2015 IEEE International Conference on Big Data), October 29, 2015, Santa Clara, CA, USA.

- [C19] **Papailiou, N., Tsoumakos, D., Konstantinou, I., Karras, P., Koziris, N.**, 2015. 'Graph-Aware, Workload-Adaptive SPARQL Query Caching'. In Proceedings of the 2015 ACM SIGMOD/PODS International Conference on Management of Data, May 31- June 4, 2015, Melbourne, Australia.
- [C20] **Doka, K., Papailiou, N., Tsoumakos, D., Mantas, C., Koziris, N.**, 2015. 'IReS: Intelligent, Multi-Engine Resource Scheduler for Big Data Analytics Workflows'. In Proceedings of the 2015 ACM SIGMOD/PODS International Conference on Management of Data (Demo Track), May 31- June 4, 2015, Melbourne, Australia.
- [C21] **Naskos, A., Stachtiri, E., Gounaris, A., Katsaros, P., Tsoumakos, D., Konstantinou, I., Sioutas, S.**, 2015. 'Dependable Horizontal Scaling Based On Probabilistic Model Checking'. In proceedings of the 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID 2015), May 4-7, 2015, Shenzhen, Guangdong, China.
- [C22] **Doka, K., Xue, M., Tsoumakos, D., Karras, P.**, 2015. 'k-Anonymization by Freeform Generalization'. In Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security (ASIACCS 2015), Singapore, April 14-17, 2015.
- [C23] **Mytilinis, I., Tsoumakos, D., Kantere, V., Nanos, A., Koziris, N.**, 2015. 'I/O Performance Modeling for Big Data Applications over Cloud Infrastructures'. In Proceedings of the 2015 IEEE International Conference on Cloud Engineering (IC2E 2015), 9-13 March, Tempe, AZ, USA.
- [C24] **Giannakopoulos, I., Tsoumakos, D., Papailiou, N., Koziris, N.**, 2015. 'PANIC: Modeling Application Performance over Virtualized Resources'. In Proceedings of the 2015 IEEE International Conference on Cloud Engineering (IC2E 2015), 9-13 March, Tempe, AZ, USA.
- [C25] **Kordoni, V., Cholakov, K., Egg, M., Way, A., Birch, A., Kermanidis, K., Sosoni, V., Tsoumakos, D., Van den Bosch, A., Hendrickx, I., Papadopoulos, M., Georgakopoulou, P., Gialama, M., Van Zaanen, M., Buliga, I., Jermol, M., Orlic, D.**, 2015. 'TraMOOC: Translation for Massive Open Online Courses'. In Proceedings of the 18th Annual Conference of the European Association for Machine Translation (EAMT2015), May 2015, European Association for Machine Translation (EAMT) Antalya, Turkey.
- [C26] **N. Nodarakis, N., Pitoura, E., Sioutas, S., Tsakalidis, A., Tsoumakos, D., Tzimas, G.**, 2014. 'Efficient Multidimensional AkNN Query Processing in the Cloud'. In Proceedings of the 25th International Conference on Database and Expert Systems Applications (DEXA 2014), Munich, Germany, September 1-5, 2014.
- [C27] **Papailiou, N., Tsoumakos, D., Konstantinou, I., Karras, P., Koziris, N.**, 2014. 'H2RDF+: An Efficient Data Management System for Big RDF Graphs'. In Proceedings of the 2014 ACM SIGMOD/PODS International Conference on Management of Data (Demo Track), Snowbird, Utah, USA.
- [C28] **Papailiou, N., Tsoumakos, D., Konstantinou, I., Karras, P., Koziris, N.**, 2014. 'Scalable Indexing and Adaptive Querying of RDF Data in the Cloud'. In Proceedings of the 6th International Workshop on Semantic Web Information Management - SWIM 2014, Snowbird, Utah, USA.
- [C29] **Giannakopoulos, I., Papailiou, N., Mantas, C., Konstantinou, I., Tsoumakos, D., Koziris, N.**, 2014. 'CELAR: Automated Application Elasticity Platform'. In proceedings of the 2014 IEEE International Conference on Big Data (BigData 2014), Washington DC, USA.
- [C30] **Tsoumakos, D., Mantas, C.**, 2013. 'The Case for Multi-engine Data Analytics'. In Proceedings of the 2013 Workshop on Middleware for HPC and Big Data Systems (MHPC'13, part of Euro-Par 2013), Aachen, Germany, August 27, 2013.
- [C31] **Papailiou, N., Konstantinou, I., Tsoumakos, D., Karras, P., Koziris, N.**, 2013. 'H2RDF+: Highperformance Distributed Joins over Large-scale RDF Graphs'. In proceedings of the 2013 IEEE International Conference on Big Data (BigData 2013), Santa Clara, CA, USA, October 6-9 2013.
- [C32] **Konstantinou, I., Tsoumakos, D., Koziris, N., Mytilinis, I.**, 2013. 'DBalancer: Distributed Load Balancing for NoSQL Data-stores'. In Proceedings of the 2013 ACM SIGMOD/PODS International Conference on Management of Data (Demo Track), New York, New York, USA, June 22-27, 2013.
- [C33] **Konstantinou, I., Kantere, V., Tsoumakos, D., Koziris, N.**, 2013. 'COCCUS: Self-configured Cost-based Query Services in the Cloud'. In Proceedings of the 2013 ACM SIGMOD/PODS International Conference on Management of Data (Demo Track), New York, New York, USA, June 22-27, 2013.
- [C34] **Tsoumakos, D., Konstantinou, I., Boumpouka, C., Sioutas, S., Koziris, N.**, 2013. 'Automated, Elastic Resource Provisioning for NoSQL Clusters Using TIRAMOLA'. In Proceedings of the 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), Delft, The Netherlands, May 13-16, 2013.
- [C35] **Konstantinou, I., Angelou, E., Tsoumakos, D., Boumpouka, C., Koziris, N., Sioutas, S.**, 2012. 'TIRAMOLA: Elastic NoSQL Provisioning through a Cloud Management Platform'. In proceedings

- of the 2012 ACM SIGMOD/PODS International Conference on Management of Data (Demo Track), Scottsdale, Arizona, USA, May 20-24, 2012.
- [C36] **Angelou, E., Papailiou, N., Konstantinou, I., Tsoumakos, D., Koziris, N.**, 2012. 'Automatic Scaling of Selective SPARQL Joins Using the TIRAMOLA System' In Proceedings of the 4th International Workshop on Semantic Web Information Management, Scottsdale, Arizona, USA, May 20, 2012
- [C37] **Papailiou, N., Konstantinou, I., Tsoumakos, D., Koziris, N.**, 2012. 'H₂RDF: Adaptive Query Processing on RDF Data in the Cloud'. In proceedings of the 21st International World Wide Web Conference (WWW2012 Demo Track), Lyon, France, April 16-20 2012.
- [C38] **Antonopoulos, P., Konstantinou, I., Tsoumakos, D., Koziris, N.**, 2012. 'Efficient Updates for Web-scale Indexes over the Cloud'. In proceedings of the Data Management in the Cloud (DMC) Workshop, Collocated with ICDE 2012, Washington, DC, April 5th 2012.
- [C39] **Sioutas, S., Panaretos, A., Karydis, I., Tsoumakos, D.**, 2012. 'SART: Dynamic P2P Query Processing in Sensor Networks with probabilistic guarantees'. In proceedings of the 27th Symposium on Applied Computing (ACM SAC 2012), March 26-30, 2012, Riva del Garda (Trento), Italy.
- [C40] **Sakkopoulos, E., Sioutas, S., Papaloukopoulos, G., Panaretos, A., Tsoumakos, D., Saltou, A.**, 2012. 'DREAM: A Distributed fRamework for customized dEployment of a vARiety of indexing engines over Million-node overlays'. In proceedings of the 27th Symposium on Applied Computing (ACM SAC 2012), March 26-30, 2012, Riva del Garda (Trento), Italy.
- [C41] **Konstantinou, I., Angelou, E., Boumpouka, C., Tsoumakos, D., Koziris, N.**, 2011. 'On the Elasticity of NoSQL Databases over Cloud Management Platforms'. In proceedings of the 20th ACM Conference on Information and Knowledge Management (CIKM 2011) – poster paper, October 24-28, 2011 – Glasgow, Scotland.
- [C42] **Doka, A., Tsoumakos, D., Koziris, N.**, 2011. 'KANIS: Preserving k-Anonymity Over Distributed Data'. In proceedings of the 5th International Workshop on Personalized Access, Profile Management, and Context Awareness in Databases (PersDB 2011), September 2, 2011 - Seattle, WA.
- [C43] **Asiki, A., Tsoumakos, D., Koziris, N.**, 2011. 'LinkedPeers: A Distributed System for Interlinking Multidimensional Data'. In proceedings of the 22nd International Conference on Database and Expert Systems Applications (DEXA 2011), Toulouse, France, August 29 - September 2, 2011.
- [C44] **Doka, A., Tsoumakos, D., Koziris, N.**, 2010. 'Brown Dwarf: A P2P Data-Warehousing System'. In proceedings of the 19th ACM International Conference on Information and Knowledge Management (demo track), Toronto, Canada, October 26-30, 2010.
- [C45] **Doka, A., Tsoumakos, D., Koziris, N.**, 2010. 'Distributing the Power of OLAP'. In proceedings of the 19th ACM International Symposium on High Performance Distributed Computing (HPDC'10), Chicago, Illinois, USA, June 23-25, 2010.
- [C46] **Doka, A., Tsoumakos, D., Koziris, N.**, 2010. 'Efficient Updates for a Shared Nothing Analytics Platform'. In Proceedings of the International Workshop on Massive Data Analytics over the Cloud (MDAC2010) in conjunction with WWW 2010, Raleigh, North Carolina, USA, April 26, 2010.
- [C47] **Konstantinou, I., Angelou, E., Tsoumakos, D., Koziris, N.**, 2010. 'Distributed Indexing of Web Scale Datasets for the Cloud'. In Proceedings of the International Workshop on Massive Data Analytics over the Cloud (MDAC2010) in conjunction with WWW 2010, Raleigh, North Carolina, USA, April 26, 2010.
- [C48] **Zisimos, A., and Doka, K., Chazapis, A., Tsoumakos, D., and Koziris, N.**, 2009. 'Optimizing Data Management in Grid Environments'. In Proceedings of the 11th International Symposium on Distributed Objects, Middleware, and Applications (DOA'09), Vilamoura, Algarve-Portugal, Nov01- 03, 2009.
- [C49] **Konstantinou, I., Tsoumakos, D., Koziris, N.**, 2009. 'Measuring the Cost of Online Load-Balancing in Distributed Range-Queryable Systems'. In Proceedings of the Ninth International Conference on Peer-to-Peer Computing, Seattle, WA, USA, 8-11 Sept. 2009.
- [C50] **Asiki, A., Tsoumakos, D., Koziris, N.**, 2009. 'An Adaptive On-line System for Efficient Processing of Hierarchical Data', In Proceedings of the 18th International Symposium on High Performance Distributed Computing (HPDC), Munich, Germany, 11-13 June 2009.
- [C51] **Doka, A., Tsoumakos, D., Koziris, N.**, 2008. 'HiPPIS: An Online P2P System for Efficient Lookups on d-Dimensional Hierarchies', In Proceedings of the 10th International Workshop on Web Information and Data Management, Napa Valley, California, 26-30 October 2008.
- [C52] **Konstantinou, I., Tsoumakos, D., Koziris, N.**, 2008. 'PASS It ON (PASSION): An Adaptive Online Load-Balancing Algorithm for Distributed Range-Query specialized Systems', In

- Proceedings of the 16th International Conference on Cooperative Information Systems (Poster Paper), Monterrey, Mexico, 12-14 November 2008.
- [C53] **Asiki, A., Doka, A., Tsoumakos, D., Koziris, N.**, 2008. 'Online Querying of Concept Hierarchies in P2P Systems', In Proceedings of the 16th International Conference on Cooperative Information Systems, Monterrey, Mexico, 12-14 November 2008.
 - [C54] **Asiki, A., Doka, A., Tsoumakos, D., Koziris, N.**, 2008. 'Support for Concept Hierarchies in DHTs', In Proceedings of the 8th International Conference on Peer-to-Peer Computing, Aachen University, Germany, 8-11 September 2008.
 - [C55] **Kantere, V., Tsoumakos, D., Sellis, T.**, 2008. 'Semantic Grouping of Social Networks in P2P Database Settings'. In proceedings of the 18th International Conference on Database and Expert Systems Applications (DEXA '07), Regensburg, Germany, 3-7 September 2007.
 - [C56] **Tsoumakos, D., Roussopoulos N.**, 2006. 'APRE: An Adaptive Replication Scheme for Unstructured Overlays'. In proceedings of the 14th International Conference on Cooperative Information Systems (CoopIS 2006), Montpellier, France, Nov 1-Nov 3 2006.
 - [C57] **Tsoumakos, D., Roussopoulos N.**, 2006. 'Analysis and Comparison of P2P Search Methods'. In proceedings of the first International Conference on Scalable Information Systems (INFOSCALE 2006), Hong Kong , May 30-June 1 2006.
 - [C58] **Tsoumakos, D., Roussopoulos N.**, 2005. 'AGNO: An Adaptive Group Communication Scheme for Unstructured P2P Networks'. In proceedings of the 2005 Euro-Par Conference, Lisbon, Portugal, August 31-September 2, 2005.
 - [C59] **Bitsakos, K., Tsoumakos, D., Roussopoulos N., Alimonos, I.**, 2005. 'A Framework for Distributed Human Tracking'. In proceedings of the 2005 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'05), Las Vegas, NV, USA, June 27-30, 2005.
 - [C60] **Kantere, V., Tsoumakos, D., Roussopoulos N.**, 2004. 'Querying Structured Data in an Unstructured P2P System'. In proceedings of the 6th ACM International Workshop on Web Information and Data Management (WIDM 2004), Washington, DC, USA, November 12-13, 2004.
 - [C61] **Tsoumakos, D., Roussopoulos N.**, 2004. 'A Framework for Sharing Voluminous Content in P2P Systems'. In proceedings of the 2004 International Conference on Parallel and Distributed Processing Techniques and Applications (PDPTA'04), Las Vegas, Nevada, USA, June 21-24, 2004.
 - [C62] **Tsoumakos, D., Roussopoulos N.**, 2003. 'Adaptive Probabilistic Search for Peer-to-Peer Networks'. In proceedings of the 3rd IEEE International Conference on P2P Computing, Linkoping, Sweden, Sept 1-3 2003.
 - [C63] **Tsoumakos, D., Roussopoulos N.**, 2003. 'A Comparison of Peer-to-Peer Search Methods'. In proceedings of the Sixth International Workshop on the Web and Databases, San Diego, CA, USA, June 12-13 2003.

0.1.3.5. Conference Proceedings

- [CP01] 'Algorithmic Aspects of Cloud Computing'. First International Workshop, ALGO CLOUD 2015, Patras, Greece, September 14-15, 2015. Revised Selected Papers
Editors: Karydis, I., Sioutas, S., Triantafillou, P., Tsoumakos, D.

0.1.3.6. Consulting - Evaluations

- [A01] Evaluating project "ΕΙΚΟΣ : ΘΕΩΡΗΤΙΚΗ ΚΑΙ ΑΛΓΟΡΙΘΜΙΚΗ ΘΕΜΕΛΙΩΣΗ ΓΙΑ ΠΡΟΣΩΠΟΚΕΝΤΡΙΚΑ ΣΥΝΕΡΓΑΤΙΚΑ ΠΛΗΡΟΦΟΡΙΑΚΑ ΣΥΣΤΗΜΑΤΑ", (under the ΘΑΛΗΣ programme), 2015.
- [A02] Swiss National Science Foundation evaluator, Project funding in Mathematics, Natural sciences and Engineering (division II), 2018

0.1.3.7. Other Publications

- [OP01] **T. Bakogiannis, I. Giannakopoulos, D. Tsoumakos and N. Koziris:** 'A Similarity-based Approach to Modeling Graph Operators'. *CoRR abs/1802.05536 (2018)*
- [OP02] **Giannakopoulos, D. Tsoumakos and N. Koziris:** 'A Decision Tree Based Approach Towards Adaptive Profiling of Distributed Applications'. *CoRR abs/1704.02855 (2017)*
- [OP03] **S. Sioutas, K. Tsihlias, A. Kosmatopoulos, A. Papadopoulos, D. Tsoumakos, K. Doka:** 'Skyline Queries in O(1) time?' *CoRR abs/1709.03949 (2017)*

0.1.4. Teaching Experience

0.1.4.1 Teaching

2010-2020 Department of Informatics, Ionian University

Teaching

Undergraduate Classes:

- [T 01] Big Data Management (since Fall 2014)
- [T 02] HY400 – Operating Systems (since Fall 2011)
- [T 03] HY110 – Introduction to Computer Programming (2010 – 2013, co-lecturer)
- [T 04] HY221 – Computer Networks II (2012-2013)
- [T 05] HY270 – Internet Technologies (since Spring 2011)
- [T 06] HY110 – Computer Programming (2011-2016, co-lecturer)
- [T 07] HY301 – Database Systems II (2018)
- [T 08] HY300 – Database Systems I (2019)
- [T 09] HY315 – Information Systems Analysis and Design (2010)

Graduate Classes:

- [T 08] Advanced Topics in Language Processing (2010, co-lecturer)
- [T 09] MSC100 – Advanced Topics in Data Base Technologies (2012, co-lecturer)
- [T 10] PMS-BIOINF103 – Programming Languages and Biological Databases (2016, co-lecturer)
- [T 11] Semantic Web Programming (2016-2018, co-lecturer)
- [T 12] MSC-DIGAPPINNOV-102 - Big Data Analytics (2019-2020, co-lecturer)
- [T 13] MSC-DDI183 - Data Processing and Analytics (2020)

2016 Department of Informatics and Telecommunications of the University of Athens

Teaching

Classes:

- [T 01] M123 - Distributed Systems (Graduate Course)

2009-2010 Computer Science Department, University of Cyprus

Teaching

Classes:

- [T 01] ΕΠΛ602 – Foundations of Internet Technologies (Graduate Course)
- [T 02] ΕΠΛ012 – Web Design Technologies
- [T 03] ΕΠΛ003 – Computer Science and Information Systems
- [T 04] ΕΠΛ132 – Programming Principles II

2000-2001 Computer Science Department, University of Maryland

Teaching Assistant

Classes:

- [TA 01] CMSC 214 – Computer Science II (C++)
- [TA 02] CMSC 424 – Data Base Systems

1999-2000 School of Electrical and Computer Engineering, National Technical University of Athens

Teaching Assistant

Classes:

- [TA 01] Data Structures

0.1.4.2. Advising of Doctoral Students

PhD Completions (co-supervision at NTUA)

Student	Title	Year
Aikaterini Doka	Distributed Database Systems	2007-2011

Ioannis Konstantinou	Distributed Systems	2007-2011
Athanasia Asiki	Distributed / Grid Systems	2007-2012
Nikolaos Papailiou	Scalable, workload-based data management for unstructured data	2011-2016
Giagkos Mytilinis	Distributed Data Management and Data Compression	2012-2019
Ioannis Giannakopoulos	Data Management and Application Profiling in Cloud Platforms	2012-2018

Dr. Tsoumakos is currently supervising the following students:

Student	Area	Year
---------	------	------

0.1.5 Professional and Research Experience

0.1.5.1 Professional Positions

0.1.5.1.1 Academic and Research Positions

2020-now	Director of the Graduate Programme: "Research Directions in Computer Science" Department of Informatics Ionian University
2017-now	Associate Professor Department of Informatics Ionian University
2012-2017	Assistant Professor Department of Informatics Ionian University
2010- 2012	Visiting Professor Department of Informatics Ionian University
2009- 2010	Visiting Professor Department of Computer Science University of Cyprus
2006- now	Postdoctoral/ Researcher position Computing Systems Lab School of Electrical and Computer Engineering National Technical University of Athens
2006- 2002	Graduate Researcher Institute for Systems Research University of Maryland

0.1.5.1.2 Consulting Services

0.1.5.1.3 Other Professional Positions

1999 - 2000	Database programmer for Energy Database School of Electrical and Computer Engineering National Technical University of Athens
1998 - 1999	Web Developer GIS Laboratory, School of Architecture National Technical University of Athens

0.1.5.2 Funded Research Projects

- 2018-2021 *HiDALGO project – HPC and Big Data Technologies for Global Systems*, H2020-INFRAEDI-2018-1. Our society faces a wide range of global challenges in different areas such as sociology, economy, ecology and technology. Solving the most vital problems in these fields often requires the analysis of a huge amount of data, and the utilization of appropriate methods and facilities for ample computations. In HiDALGO we develop a computational and data analytics environment, which enables systematic collaboration between scientists with different background and facilitates a profound understanding of several major global challenges. This will allow the decision makers to predict possible outcomes of certain social processes and act accordingly.
- 2015-2018 *TraMOOC project – Translation for Massive Open Online Courses*, H2020-644333. Massive Open Online Courses have been growing rapidly in size and impact. Yet the language barrier constitutes a major growth impediment in reaching out to all peoples and educating all citizens. TraMOOC aims at tackling this impediment by developing high-quality translation of all types of text genre included in MOOCs (e.g. assignments, tests, presentations, lecture subtitles, blog text) from English into eleven European and BRIC languages (DE, IT, PT, EL, DU, CS, BG, CR, PL, RU, ZH) that constitute strong use cases, are hard to translate into and have weak MT support, thus complying with the call objectives.
- 2014-2017 *ASAP project – An Adaptive, highly Scalable Analytics Platform*, FP7-619706, Project Technical Leader. This project proposes a unified, open-source execution framework for scalable data analytics. Data analytics tools have become essential for harnessing the power of our era's data deluge. Current technologies are restrictive, as their efficacy is usually bound to a single data and compute model, often depending on proprietary systems. The main idea behind ASAP is that no single execution model is suitable for all types of tasks and no single data model (and store) is suitable for all types of data (project Technical Leader).
- 2012-2015 *CELAR project – Automatic, multi-grained elasticity-provisioning for the Cloud*, FP7-317790, Project Technical Leader. The vision of the CELAR project is to provide automatic, multi-grained resource allocation for cloud applications. This enables the commitment of just the right amount of resources based on application demand, performance and requirements, results in optimal use of infrastructure resources and significant reductions in administrative costs (project Technical Leader).
- 2010- 2013 *ARCOMEM Project – From Collect-All Archives to Community Memories – Leveraging the Wisdom of the Crowds for Intelligent Preservation*, FP7-ICT-2009.4.1, team technical leader. The aim of this project is to leverage the wisdom of the crowds, appearing through knowledge in social media, to assist in preservation and semantic annotation of digital content.
- 2006- 2009 *GRECIA Project – Grid Enabled Access to Rich media Content*, FP6-34363, project team leader. This project aims at the design and implementation of an efficient multimedia storage, querying and management system in grid systems.
- 2002-2006 *CIP-URI ARMY Research Project: “Distributed Immune Systems for Wireless Networks”*, graduate researcher. The aim of this research project was the design and development of algorithms in order to assist in the prompt and effective detection of intrusions in Mobile Ad-hoc Networks (intrusion detection in MANETS).

0.1.6. Other

0.1.6.1. Various responsibilities in International Journals and Conferences

0.1.6.1.1. Journals

0.1.6.1.1.1. Editor in International Journals

Editor in Special issue of Frontiers (<https://www.frontiersin.org/>) "Frontiers in Big Data Time Series": <https://www.frontiersin.org/research-topics/14931/big-data-time-series>

0.1.6.1.1.2. Referee in International Journals

- [1] ACM Transactions on Internet Technology
- [2] Data and Knowledge Engineering Journal (DKE)
- [3] IEEE Transactions on Learning Technologies
- [4] Information Sciences
- [5] Journal of Computer Science and Technology (JCST)
- [6] Computer Communications (Elsevier)
- [7] IEEE Transactions on Parallel and Distributed Systems (TPDS)
- [8] Information Systems
- [9] IEEE Transactions on Knowledge and Data Engineering (TKDE)
- [10] Peer-to-Peer Networking and Applications (PPNA)
- [11] Concurrency and Computation: Practice and Experience
- [12] Computers & Electrical Engineering (Special issue on “New Trends in Humanistic Informatics: Implementations and Applications”)
- [13] Algorithms (Special Issue on "Humanistic Data Processing")
- [14] The Knowledge Engineering Review (KER)
- [15] Distributed and Parallel Databases (DAPD - Springer)
- [16] Big Data Research (BDR - Elsevier)
- [17] Computing (COMP - Springer)
- [18] International Journal of Cooperative Information Systems (IJCIS - World Scientific)
- [19] ACM Computing Surveys (ACM CSUR)
- [20] The Computer Journal (COMPJ)
- [21] Future Generation Computer Systems (FGCS - Elsevier)
- [22] Computers & Electrical Engineering (COMPELECENG - Elsevier)
- [23] IEEE Transactions on Cloud Computing (TCC-CS)
- [24] International Journal of Metadata, Semantics and Ontologies (IJMSO)
- [25] ACM Transactions on Knowledge Discovery from Data (TKDD)

0.1.6.1.2. Conferences

- Track Chair – 19th IEEE/ACM International Symposium in Cluster, Cloud, and Grid Computing (CCGrid 2019)
- PC member – International Conference on Theory and Practice of Digital Libraries (TPDL 2011)
- PC member - International Workshop on Using P2P, Grid and Agents for the Development of Content Networks (upgrade-cn2011)
- PC member – Web Information and Data Management (WIDM 2012)
- PC member – First International Workshop on Service Clouds in the Enterprise and Beyond (SCEB'12)
- PC member – IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2013, 2014, 2016, 2017, 2020)
- PC member – IEEE International Conference on Big Data 2013 (IEEE BigData 2013, 2014, 2015, 2016, 2017, 2018, 2019, 2020)
- PC member – Workshop on Middleware for HPC and Big Data Systems (MHPC'13)
- PC member – International Conference on Smart Grids and Green IT Systems (SMARTGREENS 2013, 2016, 2017, 2018, 2019, 2020)
- PC member – World Congress on Geriatrics and Neurodegenerative Disease Research (GeNeDis 2014)
- PC member – IEEE International Conference on Sustainable Computing and Communications (SustainCom 2014)
- PC member – ACM SIGMOD/PODS Conference (SIGMOD 2015)
- PC member – International Workshop on NoSQL Databases, Emerging Database Technologies and Applications (NoSQL-Net'15)
- PC member – 13th International Conference on Economics of Grids, Clouds, Systems, and Services (GECON 2016)
- PC member – Top-K Shortest Path in Large Typed RDF Graphs Challenge, ESWC Challenges, (ESWC 2016)
- PC member – The International Workshop on Semantic Big Data (SBD 2016, 2017, 2019, 2020). In conjunction with ACM SIGMOD 2016
- PC member – IEEE SERVICES 2016 - Emerging Technology Track on Elastic Systems
- PC member – International Workshop on Algorithmic Aspects of Cloud Computing (ALGO CLOUD 2016, 2017)
- PC member – International Symposium on Pervasive Systems, Algorithms, and Networks (I-SPAN 2017)
- PC member – IEEE International Conference on Data Engineering (ICDE 2017, 2018)

PC member – International Conference on Machine Learning and Data Science (ICMLDS 2017)
 PC member – ACM Conference on Information and Knowledge Management (CIKM 2017, 2018, 2019, 2020)
 PC member – IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2018)
 PC member – International Conference on Graph Computing (GC 2019, 2020)
 PC member – IEEE International Conference on Big Data Science and Engineering (BigDataSE 2020)

0.1.6.1.2.1. Scientific – Organizing Committee Member

2014 – IEEE BigData 2014 Doctoral Consortium co-Chair
 2015 – ALGO CLOUD 2015 co-Chair (collocated with Algo 2015)
 2016 – First International Workshop on Multi-Engine Data AnaLytics (MEDAL 2016) co-Chair - collocated with EDBT/ICDT 2016.
 2016 – Natural, Unconventional, and Bio-inspired Algorithms and Computation Methods Workshop (NUBACoM 2016) co-Chair – collocated with GeNeDis 2016
 2018 – 10th IEEE International Conference on Cloud Computing Technology and Science (CloudCom 2018) Workshop and Tutorial co-Chair
 2018 - Electrical Engineering STudents' European association (EESTech) Challenge Organizer
 2019 - IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid 2019) Cloud Computing track chair

0.1.6.1.2.2. Referee in International Conferences

Year	Conference	Abbrev.
2002	ACM SIGMOD International Conference on Management of Data	SIGMOD
2003	International Conference on Very Large Data Bases	VLDB
2004	International Conference on Distributed Computing Systems	ICDCS
2004	Agents and Peer-to-Peer Computing	AP2PC
2004	IEEE International Conference on P2P Computing	P2P
2005	Americas Conference on Information Systems	AMCIS
2007	ACM's Special Interest Group on Management Of Data	SIGMOD
2008	International Workshop on Data Management in Peer-to-peer systems	DAMAP
2008	International Conference on Parallel Processing	ICPP
2008	IEEE International Conference on Data Engineering	ICDE
2009	International Conference on Very Large Data Bases	VLDB
2011	International Conference on Theory and Practice of Digital Libraries	TPDL
2012	Web Information and Data Management	WIDM
2012	First Intl Workshop on Service Clouds in the Enterprise and Beyond	SCEB
2012	International Symposium on Cluster, Cloud and Grid Computing	CCGrid
2013	Workshop on Middleware for HPC and Big Data Systems	MHPC
2013	IEEE International Conference on Big Data	BigData
2014	International Conference on Smart Grids and Green IT Systems	Smartgreens
2014	International Symposium on Cluster, Cloud and Grid Computing	CCGrid
2014	World Congress on Geriatrics and Neurodegenerative Disease Research	GeNeDis
2014	International Conference on Sustainable Computing and Communications	SustainCom
2014	IEEE International Conference on Big Data	BigData
2015	ACM SIGMOD International Conference on Management of Data	SIGMOD
2015	IEEE International Conference on Big Data	BigData
2015	Workshop on NoSQL Databases, Emerging Database Technologies and Applications	NoSQL-Net
2016	International Conference on Economics of Grids, Clouds,	GECON

	Systems, and Services	
2016	IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing	CCGrid
2016	Top-K Shortest Path in Large Typed RDF Graphs Challenge	ESWC Challenges
2016	IEEE International Conference on Big Data	BigData
2016	International Workshop on Algorithmic Aspects of Cloud Computing	ALGO CLOUD
2016	International Workshop on Semantic Big Data	SBD
2016	IEEE SERVICES 2016 - Emerging Technology Track on Elastic Systems	SERVICES
2016	6th International Conference on Smart Cities and Green ICT Systems	SMARTGREENS
2017	33rd IEEE International Conference on Data Engineering	ICDE
2017	International Workshop on Semantic Big Data	SBD
2017	International Workshop on Algorithmic Aspects of Cloud Computing	ALGO CLOUD
2017	International Symposium on Pervasive Systems, Algorithms, and Networks	I-SPAN
2017	IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing	CCGRID
2017	ACM Conference on Information and Knowledge Management	CIKM
2017	International Conference on Machine Learning and Data Science	ICMLDS
2017	IEEE International Conference on Big Data	BigData
2018	34th IEEE International Conference on Data Engineering	ICDE
2018	IEEE International Conference on Big Data	BigData
2018	IEEE International Conference on Cloud Computing Technology and Science	CloudCom
2018	International Conference on Smart Cities and Green ICT Systems	SMARTGREENS
2018	ACM Conference on Information and Knowledge Management	CIKM
2019	IEEE International Conference on Big Data	BigData
2019	First International Conference on Graph Computing	GC
2019	ACM Conference on Information and Knowledge Management	CIKM
2019	The International Workshop on Semantic Big Data	SBD
2019	International Conference on Smart Cities and Green ICT Systems	SMARTGREENS
2020	International Conference on Smart Cities and Green ICT Systems	SMARTGREENS
2020	The International Workshop on Semantic Big Data	SBD
2020	IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing	CCGrid
2020	14th IEEE International Conference on Big Data Science and Engineering	BigDataSE
2020	ACM Conference on Information and Knowledge Management	CIKM
2020	IEEE International Conference on Big Data	BigData
2020	Second International Conference on Graph Computing	GC

0.1.6.3. Main Speaker in Specialized Seminars

18/03/2008	University of Crete – Foundation for Research and Technology
Multiple	University of Maryland – Comp. Sci. Dept, Database Chat
05/03/2009	University of Crete – Department of Computer Science
22/06/2009	Technical University of Crete – Dept. of Electrical and Computer Engineering
04/11/2009	University of Cyprus – Computer Science Department
23/02/2010	University of Ioannina – Computer Science Department
01/04/2011	National and Kapodistrian University of Athens – Department of Informatics
19/05/2011	National Technical Un. of Athens – Dept. of Electrical and Computer Engineering
02/03/2012	University of Cyprus – Computer Science Department
03/10/2016	University of Cyprus – Computer Science Department
08/12/2016	University of Maryland – Department of Computer Science
15/09/2017	Technical Un. of Athens – Dept. of Electrical and Computer Engineering

0.2. Citations

Up to date (30-11-2020), there have been over 2172 citations, **h-index: 19**, **g-index: 45** (Publish or Perish and Google Scholar).